Acquisition and Production of Skilled Behavior in Dynamic Decision-Making Tasks:

Semiannual Status Report*

for

NASA Research Grant NAG2-656

Principal Investigator:

Alex Kirlik
School of Industrial and Systems Engineering
Georgia Institute of Technology
Atlanta, GA 30332-0205

*As per agreement, this semiannual status report consists of two Center for Human-Machine Systems Research Technical Reports: Numbers 91-2 & 91-3.

Perceptual and Contextual Influences on Dynamic Decision-Making Performance

Alex Kirlik
Wendy J. Markert
Robert J. Shively*

Center for Human-Machine Systems Research
School of Industrial and Systems Engineering
Georgia Institute of Technology
Atlanta, GA 30332-0205

Report 91-2

Also published in *Proceedings of the 1990 IEEE International Conference on Systems, Man, and Cybernetics*, Los Angeles, CA.

*Robert J. Shively is with the U.S. Army Aeroflightdynamics Directorate, NASA Ames Research Center.

This research was supported by NASA Ames grant NAG2-656, Technical Monitor, Robert J. Shively, and an ASEE-NASA-Stanford Summer Faculty Fellowship award to Alex Kirlik.

ABSTRACT

This paper describes ongoing research investigating perceptual and contextual influences on skilled human performance in dynamic decision-making environments. The research is motivated by two general classes of findings in recent decision-making research. First, many studies suggest that the concrete context in which a task is presented has strong influences on the psychological processes used to perform the task and in subsequent performance. Second, studies of skilled behavior in a wide variety of task environments typically implicate the perceptual system as an important contributor to decision-making performance, either in its role as a mediator between the current decision context and stored knowledge, or as a mechanism capable of directly initiating activity through the development of a "trained eye." Both contextual and perceptual influences place limits on the ability of traditional utility-theoretic accounts of decision-making to guide display design, as variance in behavior due either to contextual factors or the development of perceptual skill is left unexplained. We outline a framework in which to view questions of perceptual and contextual influences on behavior and describe an experimental task and analysis technique which will be used to diagnose the possible role of perception in skilled decision-making performance.

INTRODUCTION

Two major themes run through much recent research on human decision-making. First, the psychological processes engaged when performing a given task appear to be strongly affected by the particular concrete context in which the task is presented [5,8]. Humans seem to be relatively insensitive to the formal, or deep structure of a task and tend to be strongly influenced by the problem content and contextual features of information display. Seemingly minor changes in wording or graphical display often have major consequences on resulting human behavior. Abstract problem representations (in psychology as in any science), are intended for wide range of reference and they achieve this goal by necessarily dispensing with problem content and context in favor of preserving deep structural invariance. That psychological processes may not respect such deep invariances should perhaps not come as a surprise, if we consider the dimensions along which tasks are distributed for a typical human decision-maker or problem solver. Rarely are people asked to solve a large number of problems with similar deep structure but variable content or context. Rather, the skilled decision-maker (e.g., a pilot or process control operator) works in a world of problems of similar content, presented within a fixed context, but with widely varying deep structure. There would seem to be little motivation for such a skilled performer to develop and use solution methods with greater range of reference than his or her own task environment.

This line of reasoning lends itself to interpreting a second major set of findings on human decision-making, in this case concerning skilled performance and the development of expertise. In a large number of task domains, such as chess [1], naturalististic decision-making [6], and X-ray diagnosis [7], expert performance appears to rely heavily upon an immediate perceptual evaluation of the decision situation. The expert chess player reduces the cognitive demands of elaborate search, the naturalistic decision-maker reduces the cognitive demands of enumeration and comparative evaluation of decision alternatives, and the X-ray diagnostician reduces the cognitive demands of Bayesian reasoning, all by virtue of some form of immediate perceptual processing of the problem state. The mechanisms underlying these abilities are not well understood, but it is typically hypothesized that extensive experience has enabled the decision-maker to quickly match the current situation to one that has been seen in the past, and to thereby retrieve, rather than calculate, an appropriate action to be taken. Here, perception is responsible for the pattern recognition function that encodes the situation in terms of features highly relevant to the task, to enable the activation of "chunks," "schemas," or "cases," in long-term memory.

These studies once again stress the primacy of the concrete decision context, as opposed to the abstract problem structure, in determining the psychological processes engaged in task performance. Regardless of whether a decision problem might be best formally characterized in terms of a search tree, decision tree, or through Bayesian evidential reasoning, with experience human performance appears to gravitate toward a heavy reliance upon a perceptual or recognitional mode of processing, presumably strongly determined by the perceptible information available. In addition, there is reason to believe that humans *will* perform tasks in this mode when highly skilled, despite the fact that a more formal approach to the decision task might be required to achieve optimal performance. If this is the case, human-machine system designers should attempt to ensure that the displayed information format fosters this mode of processing and allows it to function productively, rather than encourages a more analytical, but perhaps incompatible, decision strategy [4,6]. At a minimum, displays should probably be designed with this goal in mind for the many cases in which the decision problem is too complex to be attacked by formal methods or when the decision situation is characterized by unforseen contingencies.

Unfortunately, very little in the way of specific guidance is available to the designer of information displays due to our current lack of knowledge of perceptual or recognitional decision processes. One factor that makes skilled decision processes difficult to empirically identify is the power of utility-based decision processes to mimic performance produced by processes quite different those assumed by utility theory. If human decision behavior approaches optimality, a normative model of performance will be able to successfully mimic that performance since it indicates the optimal performance level. Such a model, though, only indicates that behavior has been shaped to meet task demands, and it offers little insight into the nature of the processes responsible for behavior. While this feature makes normative models valuable for certain prediction purposes [2], descriptively valid models of skilled decision-making are nevertheless required to aid system design. As Fischoff [2] has noted, for design purposes modeling "must focus on the concrete stimuli observed by operators and consider the real problems of observing, interpreting, and integrating them." (p. 279)

Our goal in this study to to attempt to diagnose and describe the possible use of a perceptual decision process in a dynamic decision-making task. Although the research discussed above has focused primarily on the pattern recognition function perception might perform to allow current and stored situations to be associated, here we do not assume that declarative memory is implicated in every perceptual decision process. The decision processes which are the focus of the current research take the form of perceptual heuristics than can be used for action selection without the need for matching presented and represented items in memory. To give an example, it appears that people use two simple perceptual heuristics (as opposed to some other form of reasoning) to judge the relative mass of two objects involved in a collision [3]. People appear to make mass judgements by assuming that faster moving objects are less massive, and that objects that ricochet backward are less massive than the object it hit. This is a case in which the identification of the perceived information used appears to provide a sufficient account of the judgement process without assuming that stored memory for collision events are consulted in the decision process. The perceptual heuristics considered here have this nature.

DIAGNOSING A PERCEPTUAL DECISION PROCESS

As with the perceptual heuristics for judging masses mentioned above, we hypothesize that skilled decision-makers sometimes use perceptual heuristics to obviate the need for elaborate reasoning. The method we propose for empirically identifying these heuristics takes advantage of the fact that, as heuristics, these processes will fail to produce optimal performance in certain decision situations. Any given display format for a decision task makes available a number of dimensions of information that can serve as the basis for a perceptual heuristic. Each dimension of displayed information, though, may be only conditionally diagnostic of the appropriate action to be taken. That is, a heuristic based on a given dimension will produce correct performance in some sub-region of the decision problem space, and sub-optimal performance in the remaining region. Given knowledge of the task constraints and payoff structure, it is possible to identify a number of possible perceptual heuristics that could be used to perform the task and associate with each an "error signature" that allows use of the heuristic to be inferred from the empirical data. These ideas will be made more precise with reference to a specific laboratory task.

Laboratory Task

The experimental task is a variant of a dynamic decision-making task originally developed by Tulga and Sheridan [9]. It should be made clear that we have made a number of changes to the task dynamics and interface in order to create an instrument that can be used to diagnose perceptual decision processes. As a result of these changes, our treatment of human performance is intended to be restricted in scope to the current version of the task only. Any extrapolations of our findings to other research using variants of this task must be made with extreme care.

The display, shown in Figure 1, shows four horizontal lines, At the start of each trial, four rectangles of variable height and width appear at the left side of the screen and begin to move to the right. Each rectangle represents a task that must be processed before it reaches its due-date (the end of the line). In the present version of the task, subjects processed a rectangle by pressing a numbered key. The keys were vertically arranged on a modified keyboard to ensure spatial compatibility with the display. In general, the subject's task is to determine an appropriate order in which to process the four rectangles and to enter the order as rapidly as possible.

Only one rectangle can be processed at a time. If the subject enters the keys 1,2,3,4 in that order, rectangle 1 will begin processing and the other rectangles will travel some variable distance across the screen during the time in which rectangle 1 is being processed. After the first rectangle has completed processing, the next rectangle will begin processing, and so on, until all the rectangles have disappeared from the screen. The subject cannot interrupt the ongoing processing of a rectangle, and may make keypresses at any time after the initiation of the trial, although a performance penalty is incurred if the subject delays between keypresses. As an inducement to rapidly determine an ordering of the rectangles, any time delay incurred between entering successive keys is reflected in a delay of the same duration between the end of the first rectangle's processing and the start of the second rectangle's processing. This feature of the task was included specifically to promote the development of rapid, possibly perceptual decision processes.

During intervals in which a rectangle is being processed, or in which processing is delayed, the remaining rectangles move across the screen at variable speeds. The speed with which a rectangle moves has been constrained to vary linearly with its area. Larger (area) rectangles move proportionately slower than smaller rectangles.

The optimal ordering is a function of a number of factors. First, the nominal value of a rectangle is indicated by its height. In calculating the actual payoff earned by processing a rectangle, though, the nominal value is divided by the distance the rectangle has traveled when it has completed processing. Therefore, a higher proportion of a rectangle's nominal value is earned the shorter the distance it has traveled. The overall payoff for a single trial is the sum of the payoffs earned for the four rectangles. Subjects are given the payoff information in qualitative form prior to the experiment, and receive a score at the end of each trial, indicating the percent of optimal performance that was achieved.

Calculating Error Signatures

The purpose of the following analysis is to identify a number of possible perceptual heuristics that could be used to order the four rectangles and to identify, for each heuristic, an expected pattern of errors. The method used is to identify an optimal ordering rule and to assess the degree to which various perceptible dimensions of information could serve as a basis for a perceptual heuristic to be used in lieu of the optimal rule. The maximum possible payoff (*P*) for each of the 24 (4!) possible orderings of the four rectangles in a single trial can be expressed as a function of the rectangle heights (*H*) and widths (*w*) as follows:

$$P = \sum_{i=1}^{4} H_i^2 \left[\frac{W_i}{\sum_{j=1}^{i} W_j} \right]$$

In this equation, *i* and *j* are indices indicating the order in which the rectangles are processed ($i = j = 1$ is the first rectangle processed, and so on). The term in brackets indicates how the delays associated with processing rectangles previous to rectangle *i* (rectangles indexed by *j*) contribute to lowering the attainable payoff for rectangle *i*, as less payoff can be earned since the rectangle will have moved across the screen. This factor indicating how attainable payoff for a rectangle is reduced by processing previous rectangles can be expressed solely in terms of the widths (processing times) of the rectangles.

We hypothesize that a perceptual decision process for performing this task will consist of ordering the four rectangles by reference to a single dimension of information on the display. Rectangle heights, widths, areas, shapes (height/width) are all dimensions that could be used to perceptually determine an ordering. There is no reason to believe that ordering the four rectangles by any single dimension will yield optimal performance in all trials. Since the ordering of any two rectangles *a* and *b* would be independent of the features of the remaining two rectangles under such a heuristic strategy, we can simply consider the degree to which the different dimensions can be used to determine optimal pairwise comparisons of rectangles. For the case of two rectangles, then, rectangle *a* should be processed prior to rectangle *b* if: (from the equation above)

$$H_a^2 + H_b^2\left[\frac{w_b}{w_a+w_b}\right] \geq H_b^2 + H_a^2\left[\frac{w_a}{w_a+w_b}\right]$$

Isolating features of individual rectangles, rectangle *a* should be processed prior to rectangle *b* if:

$$\frac{H_a^2}{w_a} \geq \frac{H_b^2}{w_b}$$

Unfortunately, this optimal pairwise ordering rule would require a comparison of rectangle features that do not seem to be readily perceptibly measurable. But while the squared height divided by the width of a rectangle does not appear to be perceptually accessible, the shapes (height divided by width) and heights alone do appear to be readily accessible to perception and mimic the optimal pairwise ordering rule to some extent. Ordering two rectangles by shape or height will, for many pairs of rectangles, result in agreement with the optimal rule, but will result in a ordering reversal for other pairs.

Diagnosing perceptual heuristics involving the displayed dimensions of height, shape, or any other dimension requires an identification of the sub-regions of the decision problem space in which use of these heuristics would result in optimal and sub-optimal performance. Figure 2 indicates the 16 rectangles that are used in the present experiment plotted in log(width) X log(height) space. Consider a pair of such rectangles of the form (7,X). The line of slope 1/2 through rectangle 7 partitions the space into those rectangles that should be ordered prior to rectangle 7 (those above the line), and those that should be ordered after rectangle 7, according to the optimal pairwise rule derived above. (The line of slope 1/2 is a graphical representation of the optimal pairwise rule derived above). The line of slope 1 through rectangle 7 partitions the space in a similar way, although this line is a graphical representation of the perceptual heuristic that orders rectangles based on shape (height/width). As can be seen from the figure, the shape heuristic mimics the optimal pairwise rule for a majority of rectangle pairs, but fails in some others (e.g., the pair (7,15). In general, the shape perceptual heuristic will fail for all pairs of the form (i,i+8) in the diagram. This region of the problem space determines the error signature that would be expected by use of the shape heuristic. Similarly, an error signature for the height heuristic (horizontal line) can be determined.

DISCUSSION

At the time of this writing we are conducting experiments using the methodology described above. Therefore, the hypotheses concerning perceptual decision processes discussed above remain tentative. Whether or not we are on the right track with the present approach, we do believe that a relatively finely grained analysis of both behavior and task environment such as the one discussed above are necessary in order to understand skilled performance. One complicating factor that makes these hypotheses difficult to test is the presence of what might be called "artifactual dimensions" of displayed information. In addition to dimensions such as rectangle height and width that have been explicitly designed to communicate meaningful information to the subject, many other contextual dimensions are also present as artifacts of the format chosen for the display. Shape is just one such artifactual dimension present on the experimental display. A real world example of an artifactual dimension might be the pattern produced by pointers on a set of analog cockpit displays. There is anecdotal evidence to suggest that crews attend to these patterns, and a move to digital state display based on a desire to improve access to individual state variable values might have the unintended and unfortunate consequence of removing this artifactual information source. We suspect that use of such sources might be relatively common in skilled performance. An example well known to the behavioral scientist is the clever subject who is able to defeat the most elegantly designed laboratory task by finding some "trick" associated with use of an artifactual information source (the whir of a disk drive or an aberration in the graphics software).

The method proposed above will also hopefully allow us to address the issue of how perceptual decision skills are acquired. Both the salience and diagnosticity of dimensions of information could be expected to influence which perceptual heuristics might be developed. It might be reasonable to assume that experience at a task has the function of moving the decision-maker from using dimensions of high initial salience but perhaps low diagnosticity, to dimensions of perhaps lower initial salience but higher diagosticity. In addition, we intend to explore how the proposed analysis might be extended to more complex task environments, in which one of a number of hierarchically organized perceptual heuristics could be variously selected by a differentiation of the decision problem space into classes of problems of specified types.

ACKNOWLEDGMENTS

This research was supported by NASA Ames grant NAG2-656, Technical Monitor, Robert J. Shively, and an ASEE-NASA Stanford Summer Faculty Fellowship award to the first author. The authors thank Neville Moray for generously providing software for this experiment, and Sandra Hart, Walter Johnson, Mary Kaiser, Dean Owen and Diane Schiano for valuable discussions.

REFERENCES

[1] Chase, W.G. and Simon, H.A. Perception in chess. *Cognitive Psychology, 4*, 1973.

[2] Fischoff, B. Modeling decision making for system design. In Elkind, Cardm Hochberg, and Huey (Eds.), *Human Performance Models for Computer-Aided Engineering*. National Academy Press, Washington D.C., 1989.

[3] Gilden, D.L. and Proffitt, D.R. Understanding collision dynamics. *Journal of Experimental Psychology: Human Perception and Performance*, Vol. 15, No. 2, 1989.

[4] Hammond, K.R., Hamm, R.M., Grassia, J., and Pearson, T. Direct comparison of the efficacy of intuitive and analytical cognition in expert judgement. *IEEE Transactions on Systems, Man, and Cybernetics, SMC-17*, No. 5, 1987.

[5] Kahneman, D. and Tversky, A. Prospect theory: An analysis of decisions under risk. *Econometrica, 47*, 1979.

[6] Klein, G. Recognition-primed decisions. In W.B. Rouse (Ed.) *Advances in Man-Machine Systems Research, Vol. 5*. JAI Press, Greenwich CT, 1989.

[7] Lesgold, A., Rubinson, H., Feltovich, P., Glaser, T., Klopfer, D. & Wang, Y. Expertise in a complex skill: Diagnosing X-ray pictures. In Chi, Glaser, and Farr (Eds.), *The Nature of Expertise*. Erlbaum, Hillsdale, NJ, 1988.

[8] Simon, H.A. and Hayes, J.R. The understanding problem: Problem isomorphs. *Cognitive Psychology, 8*, 1976.

[9] Tulga, M.K. and Sheridan, T.B. Dynamic decisions andworkload in multitask supervisory control. *IEEE Transactions on Systems, Man, and Cybernetics, SMC-10*, 1980.

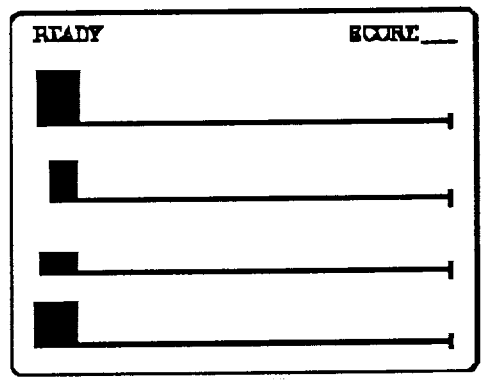

Figure 1. Experimental Display at Trial Initiation

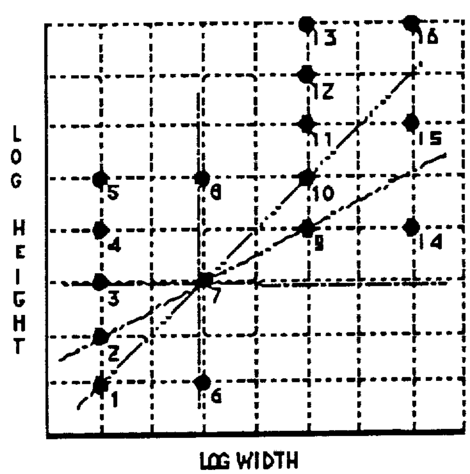

Figure 2. Decision Problem Space

Constraints on Neural Net Modeling to Support Display Design for Skilled Decision-Making

Alex Kirlik
Ling Rothrock

Center for Human-Machine Systems Research
School of Industrial and Systems Engineering
Georgia Institute of Technology
Atlanta, GA 30332-0205

Report 91-3

Also published in *Proceedings of the Second Workshop on Neural Networks: Academic / Industrial / NASA / Defense*, Auburn, AL.

This research was supported by NASA Ames grant NAG2-656, Technical Monitor, Robert J. Shively.

ABSTRACT

Designing interfaces to promote effective human decision-making requires knowledge of the psychological processes underlying skilled behavior so that information display formats can be designed to promote the development of high-level skill. Recent research suggests that skilled decision-makers often rely on perceptually intensive processes quite unlike those assumed by highly enumerative utility-theoretic decision models. What often appears to distinguish the expert from the novice is a "trained eye" for relevant information rather than the possession of elaborate decision-making procedures. Neural network modeling approaches are receiving increasing attention as candidate descriptions of perceptually-oriented decision processes since human decision behavior may gravitate toward a pattern recognition processing mode through experience. This paper discusses the suitability of neural networks as descriptions of skilled decision-making behavior, and provides a number of constraints on their design for the purpose of achieving psychologically adequate models. One major constraint concerns the specification of a psychologically appropriate input feature representation, and this issue is articulated through a discussion of an experimental study of skilled human behavior in a graphically displayed dynamic decision-making task. This experiment allowed the dimensions of graphical information used by human subjects to be identified, and therefore provides empirically motivated constraints on the design of neural network models of skilled decision behavior.

INTRODUCTION

Researchers and practitioners in human-machine systems engineering have long been interested in the problem of designing information displays to promote effective human decision-making. Candidate display formats are always motivated by the designer's theories of the psychological processes operative in human decision-making, regardless of whether these theories are explicitly formalized in a model or whether these theories consist only of implicitly held unformalized assumptions. To a large extent, the effectiveness of a display-based decision aid will be determined by the degree to which the psychological assumptions motivating the display design are consistent with the decision processes employed by the display user. To date, research on display design for decision support has almost exclusively relied upon a utility-theoretic characterization of the processes underlying decision and choice behavior. Utility theory (von Neumann and Morgernstern 1947) portrays the decision maker as exhaustively enumerating the set of decision alternatives or *actions*, estimating the probabilities of various outcomes resulting from each action, comparitively evaluating these outcomes with respect to a utility function, and finally selecting a desired action on the basis of some type of maximization operation.

Within this perspective, questions of display design for decision support concern the appropriate presentation of action attributes, probabilistic outcomes, and other information related to the use of the utility theoretic decision-making procedure. For aiding an experienced decision-maker, though, such design guidelines will be productive only to the extent that a utility-theoretic decision procedure is actually being employed by the expert decision-maker. In novel applications, these guidelines will only be effective if display designs motivated by utility theory promote the acquisition of high-level skill through experience.

Findings from the past forty years of empirical research on human decision behavior have necessitated a gradual retreat from the view that human decision-makers always behave as utility maximizers in the sense originally portrayed by utility theory. As a result, there is real concern that displays effective for supporting a utility-based decision process may provide little support for the processes skilled decision-makers actually use. To the extent that the psychological processes underlying skilled decision behavior are unlike utility-based operations, these utility-based display designs may even inhibit the development of high-level skill. As will be discussed below, it appears to be the case that the utility theoretic model may well describe decision behavior only at low and intermediate stages of skill development in a given problem domain, where high levels of skill are better described by perceptually and/or recognitionally-based processes that do not appear to rely upon the highly enumerative calculations presumed by utility theory. If this is the case, knowledge of these alternative decision processes is required to support the design of displays that are consistent with the acquisition and operation of the decision strategies characteristic of high-level skill. This paper is concerned with assessing the potential of neural network models to provide a suitable account of skilled decision processes to support the design of display formats compatible with skilled performance.

PERCEPTUAL ASPECTS OF SKILLED DECISION-MAKING

Simon (1959) observed that human information processing limitations will often prohibit a person from fully implementing a utility-theoretic decision process. Instead, Simon proposed that a person in a decision situation of realistic complexity will *satisfice* rather than optimize, in that the decision-maker will attempt to achieve an acceptable rather than optimal solution to a decision problem. The principle of satisficing may have been the first break from the extreme view of human decision-making as utility maximization, but it was by no means the last. In recent years a number of behavioral scientists have found that human decision behavior departs from normative standards in a variety of ways (e.g., see Einhorn and Hogarth (1981) and Kahneman, Slovic and Tversky (1982)), and in response theories of decision-making have been proposed that relax more and more of the the strict assumptions of utility theory. For example, Kahneman and Tversky's (1979) prospect theory not only relaxes the assumption that human estimation of probabilities will be veridical, but it also provides a *framing* stage that allows contextual aspects of a decision problem to influence decision behavior, although these contextual features have no representation in the normative utility-theoretic model.

More recently, a number of researchers studying skilled decision behavior in environments of high complexity have documented an additional set of inadequacies of the utility-based approach. Unlike the research referenced above that maintains the overall picture of decision-making as (albeit imperfect) option enumeration, evaluation and selection, these new findings suggest a radically different account of the processes underlying skilled behavior. Klein's (1989) studies of skilled decision-making in naturalistic contexts has led him to propose that highly experienced decision-makers rarely undergo a comparative evaluation process while making decisions, but rather select a promising action based on a rapid perceptual assessment of the decision environment. In both the combinatorially complex game of chess (Chase and Simon 1973) and the challenging task of X-ray diagnosis (Lesgold et al. 1988), highly experienced experts appear to rely upon a rapid perceptual evaluation of the decision situation as the basis for action selection.

Such skills are often thought to rely upon the activation of memories for previous decision problems, and in this case perception is presumed to be responsible for the pattern recognition function that allows current and stored experiences to be associated. This recognitional as opposed to purely perceptual view, however, has only limited empirical support (but see Logan 1988), for it is not yet known whether the ability to summon memories for previous encounters is a critical component of skilled decision-making behavior. Our own previous research on skilled performance using graphically displays, for example, is based on the view that skilled decision behavior can be generated by a tuning of the perceptual system to highly relevant features of the environment to allow the direct perceptual selection of action without the need to assume memories are consulted in the decision-making process, in the tasks of route planning (Kirlik, Miller and Jagacinski 1988), supervisory control (Kirlik, Miller and Jagacinski 1989), and dynamic decision-making (Kirlik, Markert and Shively 1990). In the following, we shall be concerned mainly with skilled decision processes of this purely perceptual, rather than the more complex perceptual + recognitional, type.

A NEURAL NETWORK ALTERNATIVE

The possibility that skilled decision behavior is heavily reliant upon a perceptually intensive evaluation of the environment suggests that neural network models may have the potential to model the acquisition and production of decision-making skills. One central reason that neural network models may be attractive in this respect is that they appear to offer a candidate description of how an expert may develop a "trained eye" through experience at a decision task, and how this capability might obviate the need for a more cognitively elaborate decision process to achieve successful performance. The ability of adaptive networks to become attuned to the most highly diagnostic features of a decision problem may shed some light on the principles of perceptual development underlying the acquisition of expertise, and may therefore be illuminating as to the way in which information displays should be designed to promote skill acquisition and skilled decision-making performance.

Unfortunately, there is still much work to be done before it can be determined whether neural network models can provide an empirically adequate model of human decision-making skills. As Pavel, Gluck and Henkle (1988) have noted, most current neural network architectures and learning techniques are simply too unconstrained to mimic human learning and psychologically appropriate constraints are not yet known. In that work, for example, the authors investigated adding a *minimal congfiguration* constraint to a standard three-layer network in an attempt to mimic human generalization behavior, hypothesizing that a network with a minimal number of hidden units would produce a psychologically adequate model. The minimal configuration constraint did not produce adequate generalization behavior, however, and the authors have suggested exploring an additional set of constraints dealing with initial conditions of connection weights.

We believe that an identification of important constraints on network design can proceed most efficiently by detailed empirical investigations of the aquisition and use of skilled human decision processes. Thus, we advocate a primarily bottom-up modeling orientation that will hopefully avoid many false starts in proposing perhaps elegant, yet psychologically irrelevant, modeling architectures and learning rules. Other, more top-down approaches may certainly be more appropriate for achieving engineering solutions to various tasks, but the demands of information display design in human-machine systems requires a model that faithfully represents the nature of the processing activities underlying skilled human behavior. In this spirit, we present below a discussion of some issues that must be considered in constructing an empirically adequate neural network model of perceptually-oriented decision processes.

Input Feature Representations

One primary motivation for considering a neural network modeling approach is to address the issue of how an expert decision-maker may become attuned to the most highly diagnostic dimensions of information in a given task environment. As Dawes and Corrigan (1974) have noted, experts in a particular decision environment seem to be much better at selecting and coding the relevant input variables, or features, than they are at integrating feature information after it has been selected. This empirical phenomenon is one major reason that simple linear regression models are typically quite successful in out-performing even expert decision-making behavior (Goldberg 1968; Johnson 1988). These simple models perform so well at various judgemental tasks (e.g., medical and psychological diagnosis) because considerable human expertise has gone into selecting the independent variables that are relevant to the decision problem. After the relevant features have been selected, it may not be surprising that these models achieve better performance than human experts since they combine information in an optimal (i.e., least squares) fashion, whereas human abilities to perform information integration are quite limited (Dawes, 1979).

Since much of the responsibility for human decision-making expertise may be due to this feature selection capability, it is highly desirable to understand the psychological constraints that determine which input features will be attended to by the human expert. From a modeling perspective, this issue is directly related to the problem of specifying the primitive input feature representation used for environmental description. The reason that the choice of environmental description is so critical is that the mere act of describing the environment strongly constrains the possible behaviors that a model using such a description can exhibit, and also partially determines the information integration activites necessary to perform any classification task. Any environmental features that are equivalently classed, or else abstracted away entirely in the environmental description are rendered incapable of producing behavioral variance. On the other hand, the environmental description may be too finely rather than too coarsely grained (with respect to the input features used by the human), with the result that significant amounts of information integration might be necessary to reconstruct within the model what are actually primitive environmental features from the perspective of the human expert. The danger here is that the model will suggest that a good deal of possibly cognitive-level information integration is required to obtain supposedly "high-order" features, whereas the human expert may be actually perceptually detecting a high-order environmental feature directly.

The problem of environmental representation is by no means new, but it is especially critical for the purpose of psychological modeling with neural networks. For example, Pao (1989) has noted that innappropriate choices for a network's primitive feature representation can lead to the need for very complex decision or categorization rules, wheras incisive choices can lead to much simpler decision rules to perform the same decision task. Schuurmans and Schaeffer (1989) have made a similar point in the context of developing classsifier system models based on genetic algorithms, where they note that the selection of primitive input features for such models has the potential to either create or avoid a combinatorial explosion in the search for rules during the learning process. The ability of Tesauro and Sejnowski's (1989) neural network to successfully learn to play the game of backgammon was highly dependent on the choice of input features used by human experts in that domain (e.g., "slotting" and "stripping"). As the arguments above suggest, much of the "intelligence" involved in skilled performance may be in the selection of these features for primitive environmental representation, rather than in the way they are internally processed, and therefore the fact that one must already be a domain expert to construct a successful learning model (Tesauro 1989; Schuurmans and Schaeffer 1989) severely compromises the ability of such approaches to provide psychological insights into the nature of the development of expertise.

Input representation problems are particularly acute when considering human decision-making behavior in the context of a graphical information format, such as on a situation or map display (Kirlik 1989). In map displays, for example, it may be the case that supposedly high-order relations among map features are actually more easily perceptually available than are the first-order features themselves. Consider the simple problem of attempting to determine a path for a vehicle between two waypoints on a map. If, say, distance and direction between the waypoints are important information sources, note that a Cartesian representation of the map display space will require that information integration be performed to obtain these "high-order" features from the first-order features (the coodinates of the waypoints), most likely using the properties of the right triangle. If, on the other hand, a polar representation of the display space was used (with origin at the point of foveation), the distance and direction between two waypoints could be obtained directly by foveating at one waypoint and simply reading the (radius, angle) coordinates of the other waypoint directly from the environmental representation. Note how the choice of primitive feature representation for the map display strongly determinines the complexity of the processes that would be needed to model skilled human performance in this task, and it could be argued that the potential benefits of graphical displays as decision-aids cannot be understood until appropriate display representations are found. In a later section of this paper, we will discuss some empirical and analytical approaches we are currently using in order to better understand the constraints on input feature selection by human experts, in the hope that a principled approach to input feature representation for decision-making tasks will be forthcoming.

Feature Salience

Even supposing a psychologically adequate environmental representation can be specified, it is quite likely that that the salience of various input features strongly influences the dimensions of information used by a human decision-maker in a particular task environment. For example, Gilden and Proffitt (1989) have demonstrated the importance of information salience in the processes underlying human judgments about simple physical events. Through training, neural networks have the capability of detecting the diagnosticity of an input feature for performing a particular task, but there is no simple provision allowing networks to be sensitive the the different degrees of psychological salience that various dimensions of input information may possess. In Kirlik, Markert and Shively (1990) we advanced the hypothesis that through experience a human decision-maker may move from using dimensions of high initial salience but perhaps low diagnosticity, to the use of dimensions with high diagnosticity but perhaps low initial salience. Initially, then, we suggest that decision-makers may be biased to use dimensions of information that are easy to perceptually measure, even though these dimensions may not be highly informative. Only through experience may a human decision-maker learn to attend to dimensions of information of very low initial salience if these dimensions are highly informative. In a later section we describe experimental work underway to assess the validity of these hypotheses.

The problem of describing sensivity to salience places two constraints on the design of a psychologically plausible neural network model. First, provisions must be made to explicitly represent salience, either in the environmental representation or in the model. Although it is normal to speak of salience as a property of environmental information, the fact that the salience of a dimension of information may change through experience with the development of perceptual sensitivity suggests that salience might best be described as a dynamic property of the perceptual response. What is suggested, then, is that to appropriately achieve attention allocation to various input features these features must be described not only by their currently perceived diagnosticity (as reflected by connection weights eminating from the feature nodes), but also by their current perceptual salience. A second constraint on network design concerning salience is the specification of mechanisms for the dynamic updating of salience values, as these values are changed through the development of perceptual sensitivity. As a large component in information display design is the manipuation of information salience (e.g., though color, blinking, etc.), models of display-aided decision processes must produce behavioral variance to salience so that they can be used in a predictive process of display evaluation.

The End Product of Learning

Finally, we note that since a skilled decision-maker's ability to select diagnostic dimensions of information is probably superior to his or her ability to integrate information from multiple sources, it is quite likely that humans may have a tendency to use relatively few of the potentially large number of available sources of information present in a given environment. Provided that a small number of dimensions of information can be used to obtain acceptable (but perhaps suboptimal) performance, the decision-maker may indeed process only a relatively small set of the information available in order to make a decision. It has been demonstrated, for example, that increasing the amount of information available to a decision-maker may increase the decision-maker's confidence in his or her judgements, but it may not increase the accuracy of those judgements (Slovic and Lichtenstein 1971).

These observations suggest that psychologically appropriate neural network models must be constrained to use only a relatively small subset of the available input features when fully trained. The optimal solution of a given decision problem may indeed require attention to a large number of input features, but redundancy among input feature values and/or low diagnosticity values of certain features may be exploited to allow reasonable decision accuracy with attention to a much smaller set of input features. As will be described below, we are exploring the use of the network*skeletonization* process described by Mozer and Smolensky (1989) as a potential method of pruning the set of input features to which a neural network model will attend at asymptotic performance.

AN EMPIRICAL APPROACH

We are currently using an empirical approach to identify constraints on appropriate models of skilled decision-making behavior. The laboratory task described below is being used to investigate the hypothesis that skilled decision-makers sometimes use perceptual decision processes as heuristics to obviate the need for elaborate reasoning. From a modeling perspective, our interest is in the potential of neural networks to mimic the rate and nature of the aquisition of these perceptual decision processes. The method are using to empirically identify these decision processes takes advantage of the fact that as heuristics these perceptual processes will fail to produce optimal performance in certain decision situations. The pattern of errors generated by human subjects can therefore be examined to infer the nature of the perceptual heuristics used, and can therefore determine to which dimensions of information these decision-makers attended during skilled performance. Identifying the nature of information used by these decision-makers provides a rich source of constraint on the design of appropriate neural network models for mimicking human learning and skilled decision-making behavior.

Laboratory Task

The experimental task is a variant of a dynamic decision-making task originally developed by Tulga and Sheridan (1980). A brief description of the task is given below and a more extended treatment can be found in Kirlik, Markert and Shively (1990). The display, shown in Figure 1, shows four horizontal lines, At the start of each trial, four rectangles of variable height and width appear ate keys 1,2,3,4 in that order, rectangle 1 will begin processing and the other rectangles will travel some variable distance across the screen during the time in which rectangle 1 is being processed. After the first rectangle has completed processing, the next rectangle will begin processing, and so on, until all the rectangles have disappeared from the screen. The subject cannot interrupt the ongoing processing of a rectangle, and may make keypresses at any time after the initiation of the trial, although a performance penalty is incurred if the subject delays between keypresses. As an inducement to rapidly determine an ordering of the rectangles, any time delay incurred between entering successive keys is reflected in a delay of the same duration between the end of the first rectangle's processing and the start of the second rectangle's processing. This feature of the task was included specifically to promote the development of rapid, possibly perceptual decision processes.

During intervals in which a rectangle is being processed, or in which processing is delayed, the remaining rectangles move across the screen at variable speeds. The speed with which a rectangle moves has been constrained to vary linearly with its area. Larger (area) rectangles move proportionately slower than smaller rectangles. The task is somewhat complex as the optimal ordering of the four rectangles is a function of a number of factors. First, the nominal value of a rectangle is indicated by its height. In calculating the actual payoff earned by processing a rectangle, though, the nominal value is divided by the distance the rectangle has traveled when it has completed processing. Therefore, a higher proportion of a rectangle's nominal value is earned the shorter the distance it has traveled. The overall payoff for a single trial is the sum of the payoffs earned for the four rectangles. Subjects are given the payoff information in qualitative form prior to the experiment, and receive a score at the end of each trial, indicating the percent of optimal performance that was achieved. Each experimental session consisted of a series of 80 trials, and subjects were run for a total of ten sessions over a ten day period.

Calculating Error Signatures

The purpose of the following analysis is to identify a number of possible perceptual heuristics that could be used to order the four rectangles and to identify, for each heuristic, an expected pattern of errors. The method used is to identify an optimal ordering rule and to assess the degree to which various perceptible dimensions of information could serve as a basis for a perceptual heuristic to be used in lieu of the optimal rule. The maximum possible payoff (P) for each of the 24 (4!) possible orderings of the four rectangles in a single trial can be expressed as a function of the rectangle heights (H) and widths (w) as follows:

$$P = \sum_{i=1}^{4} H_i^2 \left[\frac{W_i}{\sum_{j=1}^{i} W_j} \right]$$

In this equation, i and j are indices indicating the order in which the rectangles are processed ($i = j = 1$ is the first rectangle processed, and so on). The term in brackets indicates how the delays associated with processing rectangles previous to rectangle i (rectangles indexed by j) contribute to lowering the attainable payoff for rectangle i, as less payoff can be earned since the rectangle will have moved across the screen. This factor indicating how attainable payoff for a rectangle is reduced by processing previous rectangles can be expressed solely in terms of the widths (processing times) of the rectangles.

We hypothesize that a perceptual decision process for performing this task will consist of ordering the four rectangles by reference to a single dimension of information on the display. Rectangle heights, widths, areas, shapes (height/width) are all dimensions that could be used to perceptually determine an ordering. There is no reason to believe that ordering the four rectangles by any single dimension will yield optimal performance in all trials. Since the ordering of any two rectangles a and b would be independent of the features of the remaining two rectangles under such a heuristic strategy, we can simply consider the degree to which the different dimensions can be used to determine optimal pairwise comparisons of rectangles. For the case of two rectangles, then, rectangle a should be processed prior to rectangle b if: (from the equation above)

$$H_a^2 + H_b^2 \left[\frac{w_b}{w_a + w_b} \right] \geq H_b^2 + H_a^2 \left[\frac{w_a}{w_a + w_b} \right]$$

Isolating features of individual rectangles, rectangle a should be processed prior to rectangle b if:

$$\frac{H_a^2}{w_a} \geq \frac{H_b^2}{w_b}$$

Unfortunately, this optimal pairwise ordering rule would require a comparison of rectangle features that do not seem to be readily perceptibly measurable. But while the squared height divided by the width of a rectangle does not appear to be perceptually accessible, the shapes (height divided by width) and heights alone do appear to be readily accessible to perception and mimic the optimal pairwise ordering rule to some extent. Ordering two rectangles by shape or height will, for many pairs of rectangles, result in agreement with the optimal rule, but will result in a ordering reversal for other pairs.

Diagnosing perceptual heuristics involving the displayed dimensions of height, shape, or any other dimension requires an identification of the sub-regions of the decision problem space in which use of these heuristics would result in optimal and sub-optimal performance. Figure 2 indicates the 16 rectangles that are used in the present experiment plotted in log(width) X log(height) space. Consider a pair of such rectangles of the form (7,X) in the figure. The line of slope 1/2 through rectangle 7 partitions the space into those rectangles that should be ordered prior to rectangle 7 (those above the line), and those that should be ordered after rectangle 7, according to the optimal pairwise rule derived above. (The line of slope 1/2 is a graphical representation of the optimal pairwise rule derived above). The line of slope 1 through rectangle 7 partitions the space in a similar way, although this line is a graphical representation of the perceptual heuristic that orders rectangles based on shape (height/width). As can be seen from the figure, the shape heuristic mimics the optimal pairwise rule for a majority of rectangle pairs, but fails in some others (e.g., the pair (7,15). In general, the shape perceptual heuristic will fail for all pairs of the form (i,i+8) in the diagram. This region of the problem space determines the error signature that would be expected by use of the shape heuristic. Similarly, an error signature for the height heuristic (horizontal line) can be determined.

INSERT FIGURE 2 HERE

Note that with this analysis of the decision problem space, the use of each possible perceptual heuristic can potentially be inferred from the data produced by each laboratory subject. In this way, this analysis helps us identify the nature of the information used by subjects as they acquire decision skills in this task, with the result that certain constraints on appropriate neural network model design can be empirically determined. A minimal necessary condition for a neural network model of human behavior in this task is that it learns to use the same dimensions of displayed information as do subjects, and the methdology described above will allow us to validate our models with respect to this critical constraint.

Preliminary Results

We are currently performing detailed analysis of data produced by five subjects using the methodology described above. At this point, it appears that the behavior of two of the five subjects is well described as the stable use of a single dimension of displayed information for ordering the four rectangles over the last five sessions of training. The error profile produced by one subject is highly similar to the error signature predicted by examination of the height-ordering perceptual heuristic, while the error profile generated by another subject is highly similar to the error signature expected through use of the area-ordering heuristic. The other three subjects appeared to have used a mixture of the possible heuristics, or else a heuristic based on a dimension of perceptual information not yet identified. Rigorous statistical tests of these conclusions, as well as detailed analyses of the other three subjects are not available at this time.

One interesting observation that can be made even from the preliminary results of this experiment is the use of what might be called "artifactual dimensions" of graphically displayed information. One subject appeared to order the rectangles by a simple perceptual comparison of the areas of the four rectangles. It is important to note that the laboratory display was not intentionally designed to communicate any meaningful task information by the dimension of rectangle area. Rather, this information source is an artifact of display design, as it might not have been present if a different display format had been used. Rectangle shape is another such artifactual information source present on the display, and our results appear to suggest that subjects may learn to use such artifactual information sources if these sources are perceptually salient and somewhat diagnostic, even though they were not explicitly designed to communicate useful information. We suspect that this phenomenon is not restricted to this laboratory task environment, as in many cases graphical displays make perceptually salient high-order relations among task variables, and although the potential benefit of such relational information may not be predicted by the display designer, it may nevertheless be discovered by the highly skilled display user.

CONSTRAINTS ON NEURAL NET MODELS

The previous discussion of modeling constraints motivated by a review of the literature, combined with the emerging results of our own experimental work, suggest a number of constraints on the design of a psychologically adequate neural network model of learning and performance in this laboratory task. We suggest that the following constraints may also be applicable to modeling skilled decision-making processes in other perceptually rich (i.e., non-verbal) problem domains.

Relational Input Features

Input features must be determined by a finely grained examination of the perceptually available displayed information. What may be thought to be high-order relational features may actually be primitive single input features from the standpoint of perceptual sensitivity. In the present experiment, rectangle area, for example, may be just as easily measurable as rectangle height and width, and it should therefore be included in the input feature set used by a neural net model. See also Gluck, Bower and Hee (1989) for a demonstration of the use of of configural input features in neural net modeling.

Artifactual Input Features

Simply because various dimensions of graphically displayed information have not been intentionally designed to communicate meaningful information does not mean that these dimensions will always be ignored by display users. For modeling a decision-maker in a given task domain, one cannot use a utility-theoretic model, for example, to exhaustively determine network input features, since many other dimensions of information are likely to be available but have no representation in the constructs of the utility-based model. If other than utility-based processes are being used, other than utility-based information is likely to be used.

Input Feature Salience

Subjects in our experimental task had available a dimension of information that could have led to optimal pairwise orderings of rectangles, although this dimension of information was not initially salient and could therefore probably not be used. At least one, and perhaps more, subjects appeared to rely upon rectangle height, which is a very salient dimension of information although not perfectly diagnostic. We suggest that variance in human behavior will be generated as a function of both salience and diagnosticity, although to our knowledge salience has not been well represented in neural network models to this point. Due to the problems involved in attempting to measure salience, we expect this constraint on neural network modeling to present challenges for a long time to come.

Dynamic Feature Salience

The skilled decision-maker may develop perceptual sensitivity to dimensions of information of very low initial salience (e.g., as in wine tasting, chess, and music appreciation). Although the duration of training and/or the abstract nature of the information display in the present experiment may not provide for the development of perceptual sensitivity in this task, provisions for dynamic measures of input feature salience are likely to be necessary for modeling skilled decision-making in naturalistic contexts.

Size of Attended Feature Set

A skilled human decision-maker may only attend to a relatively small subset of the information available in a given task environment. We are currently exploring the skeletonization technique developed by Mozer and Smolensky to prune lowly diagnostic nodes from input and hidden layers, focusing the network on the most highly diagnostic information. Without focusing techniques such as this, the nature of many standard learning rules such as back propogation may distribute significant connection weight values widely thoughout a network, as there is no constraint on such learning rules that would trade off processing complexity for small losses in decision accuracy, as might human decision-makers.

CONCLUSIONS

Based on a review of a variety of findings in decision-making research and our own laboratory work in progress we have presented a number of constraints on the design of psychologically plausible neural network models of skilled decision behavior. Naturally, a large number of other (e.g., physiological) constraints might be suggested as well, but those outlined above capture what we believe to be many critical aspects of skilled human decision-making. Whether or not we are on the right track with the present approach, we feel that appropriately constrained neural networks might provide an opportunity for describing perceptual aspects of skilled decision behavior, and might provide an important tool for the design of displays to promote the acquistion of decision-making skills.

ACKNOWLEDGMENTS

This research was supported by NASA Ames grant NAG2-656, technical monitor Robert J. Shively. The authors thank Neville Moray for providing software necessary for the laboratory task, and Wendy Markert and Merrick Kossack for their assistance with the experiment and data analysis.

REFERENCES

Chase, W.G. and H.A. Simon. 1973. "The mind's eye in chess." In *Visual Information Processing*, W.G. Chase, Ed. Academic Press, New York, NY.

Dawes, R.M. 1979. "The robust beauty of improper linear models in decision making." *American Psychologist, 34*: 571-582.

Dawes, R.M. and B. Corrigan. 1974. "Linear models in decision making." *Psychological Bulletin, 81*: 95-106.

Einhorn, H.J. and R.M. Hogarth. 1981. "Behavioral decision theory: Processes of judgement and choice." *Annual Review of Psychology, 32*: 53-88.

Gilden, D.L. and D.R. Proffitt. 1989. "Understanding Collision Dynamics." *Journal of Experimental Psychology: Human Perception and Performance.* Vol. 15, No. 2.

Goldberg, L.R. 1968. "Simple models or simple processes? Some research on clinical judgements." *American Psychologist, 23*: 483-496.

Gluck, M.A., Bower, G.H. and M.R. Hee. 1989. "A Configural-Cue Network Model of Animal and Human Associative Learning." *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society.* Ann Arbor, MI.

Johnson, E.J. 1988. "Expertise and decision under uncertainty: Performance and process." In *The Nature of Expertise*, M.T.H. Chi, R. Glaser and M.J. Farr, Eds. Erlbaum, Hillsdale, N.J.

Kahneman, D. and A. Tversky. 1979. "Prospect theory: An analysis of decision under risk." *Econometrica, 47*: 263-291.

Kirlik, A., Miller, R.A. and R.J. Jagacinski. 1988. "Route planning as a perceptual task." *Proceedings of the 23rd Annual Conference on Manual Control*, Cambridge, MA.

Kirlik, A., Miller, R.A. and R.J. Jagacinski. 1989. "A process model of skilled human performance in a dynamic uncertain environment." *Proceedings of the 1989 IEEE International Conference on Systems, Man, and Cybernetics*, Cambridge, MA.

Kirlik, A. 1989. "Skilled Human Interaction with Graphical Displays: Perceiving Designed Representations." Paper presented at the 1989 ACM SIGCHI Workshop on Computer-Human Interaction in Aerospace Systems. Seattle, WA.

Kirlik, A., Markert, W.J., and R.J. Shively. 1990. "Perceptual and contextual influences on dynamic decision-making performance." *Proceedings of the 1990 IEEE International Conference on Systems, Man, and Cybernetics*, Los Angeles, CA.

Klein, G. 1989. "Recognition-primed decisions." In *Advances in Man-Machine Systems Research, Vol. 5*, W.B. Rouse, Ed. JAI Press, Greenwich, CT.

Lesgold, A., Rubinson, H., Feltovich, P., Glaser, T., Klopfer, D. and Y. Wang. "Expertise in a complex skill: Diagnosing x-ray pictures. In *The Nature of Expertise*, M.T.H. Chi, R. Glaser and M.J. Farr, Eds. Erlbaum, Hillsdale, N.J.

Logan, G.D. 1988. "Toward an instance theory of automatization." *Psychological Review*, Vol. 95, No. 4: 492-527.

Mozer, M.C. and P. Smolensky. 1989. "Skeletonization: A technique for trimming the fat from a network via relevance assessment." In *Advances in Neural Information Processing Systems 1*, D.S. Touretzky, Ed. Morgan Kaufmann Publishers, Palo Alto, CA.

Pao, Y. 1989. *Adaptive Pattern Recognition and Neural Networks.* Addison-Wesley, New York, NY.

Pavel, M., Gluck, M.A. and V. Henkle. 1988. "Constraints on adaptive networks for modeling human generalization." *Proceedings of the November 1988 Neural Information Processing Systems (NIPS) Conference; Denver, CO.* Morgan Kaufman, San Mateo, CA.

Schuurmans, D. and J. Schaeffer. 1989. "Representational difficulties with classifier systems." *Proceedings of the Third International Conference on Genetic Algorithms:* 328-333.

Simon, H.A. 1959. "Theories of decision-making in economics and behavioral science." *American Economic Review, 49*: 253-280.

Slovic, P. amd S. Lichtenstein. 1971. "Comparison of Bayesian and regression approaches to the study of information processing in judgement. *Organizational Behavior and Human Performance, 6*: 649-744.

Tesauro, G. and T.J. Sejnowski. 1989. "A parallel network that learns to play backgammon." *Artificial Intelligence*, in press.

Tulga, M.K. and T.B. Sheridan. 1980. "Dynamic decisions and workload in multitask supervisory control." *IEEE Transactions on Systems, Man, and Cybernetics, SMC-10.*

von Neumann, J. and O. Morgernstern. 1947. *Theory of Games and Economic Behavior* (2nd Ed.). Princeton University Press, Princeton, N.J.

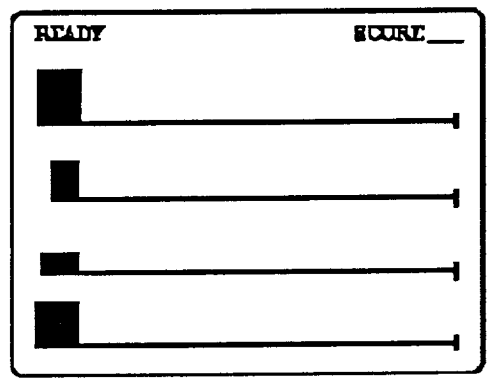

Figure 1. Experimental display at trial initiation

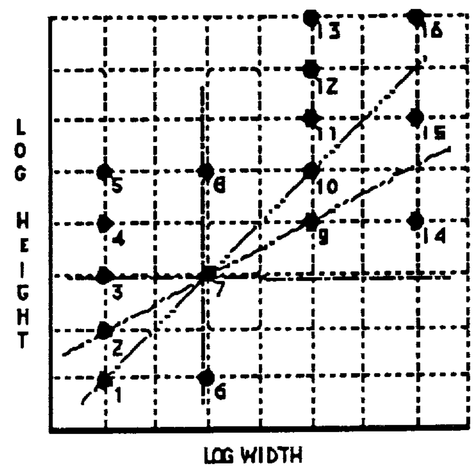

Figure 2. Decision problem space showing the 16 rectangles used as stimuli. The line of slope 1/2 is the optimal pairwise ordering rule. Various perceptual heuristics are shown by the other lines, each based on a different dimension of graphically displayed information. See text for explanation.